DRAL

*Daresbury Laboratory
Rutherford Appleton Laboratory*

RAL Report
RAL-95-029

Element Resequencing for use with a Multiple Front Algorithm

J A Scott

SW 3520

March 1995

Rutherford Appleton Laboratory Chilton DIDCOT Oxfordshire OX11 0QX

**DRAL is part of the Engineering and Physical
Sciences Research Council**

The Engineering and Physical Sciences Research Council
does not accept any responsibility for loss or damage arising
from the use of information contained in any of its reports or
in any communication about its tests or investigations

Element resequencing for use with a multiple front algorithm

J. A. Scott

ABSTRACT

The multiple front algorithm is an extension of the frontal method to allow parallelism to be exploited in the solution process. The finite-element domain is partitioned into a number of subdomains and a frontal decomposition is performed on each subdomain separately. For a given partitioning of the domain, the efficiency of the multiple front algorithm depends on the ordering of the elements within each subdomain. We look at the limitations of existing element reordering algorithms when applied to a subdomain and consider how these limitations may be overcome. Extensive numerical experiments are performed on a range of practical problems and, on the basis of the results, we propose a new element resequencing algorithm for use with a multiple front algorithm.

Keywords : sparse matrices, frontal methods, Gaussian elimination, finite-element equations.

AMS(MOS) subject classification : 65F05, 65F50.

CR classification system : G.1.3[Numerical Linear Algebra]: linear systems: sparse systems.

Computing and Information Systems Department,
Atlas Centre,
Rutherford Appleton Laboratory,
Oxon OX11 0QX.

May 1995.

CONTENTS

1	Introduction.....	1
2	Graphs and finite-element domains	3
3	Element reordering using MC43	5
4	Reordering elements in a subdomain	8
	4.1 Approach I.....	8
	4.2 Approach II	9
	4.3 Element relabelling	10
	4.4 Selecting the element order	11
	4.5 Coping with more than one component	12
5	The test problems	13
6	Conclusions.....	23

1 Introduction

In this paper, we are concerned with the solution of $n \times n$ linear systems of equations

$$\mathbf{Ax} = \mathbf{b}, \quad (1.1)$$

where $\mathbf{A}$ is a large sparse matrix arising from finite-element analysis. The matrix $\mathbf{A}$ is a sum of elemental matrices

$$\mathbf{A} = \sum_{l=1}^m \mathbf{A}^{[l]}. \quad (1.2)$$

Each matrix $\mathbf{A}^{[l]}$ has entries only in the principal submatrix corresponding to the variables in element l and represents contributions from this element. This principal submatrix is assumed to be dense. The matrix $\mathbf{A}$ may be unsymmetric but the form (1.2) implies that the sparsity pattern is symmetric with nonzero diagonal entries. One possible direct solution method for (1.1), and the one which is still frequently the method of choice in many structural engineering applications, is the frontal method (see, for example, Irons 1970, Hood 1976, Duff 1983).

The efficiency of a frontal scheme, in terms of both storage and computation time, is dependent upon the ordering of the elements. This is because, in the frontal method, the system matrix $\mathbf{A}$ is never assembled explicitly but the assembly and Gaussian elimination processes are interleaved, with each variable being eliminated as soon as its row and column are fully summed, that is, after its last occurrence in a matrix $\mathbf{A}^{[l]}$. This allows all intermediate working to be performed in a full matrix, termed the *frontal matrix*, whose rows and columns correspond to variables that have not yet been eliminated but occur in at least one of the elements that have been assembled. Since the order of the frontal matrix increases when a variable appears for the first time and decreases whenever a variable is eliminated, the order in which the elements is input is critical.

Let us introduce some notation. With reference to equation (1.1), column j is said to be *active* in row i if $j \geq i$ and there is a nonzero entry in column j with a row index $k \leq i$. Letting f_i denote the number of active columns in row i , the *maximum wavefront* of the matrix $\mathbf{A} = (a_{ij})$ is given by

$$F = \max_{1 \leq i \leq n} \{f_i\}. \quad (1.3)$$

The *root-mean-squared* (r.m.s.) *wavefront* is defined to be

$$\tilde{F} = \left(\frac{1}{n} \sum_{i=1}^n f_i^2 \right)^{1/2}. \quad (1.4)$$

The *profile* of the matrix $\mathbf{A}$ is the total number of coefficients in the lower triangle when any zero ahead of the first entry in its row is excluded. That is,

$$P = \sum_{i=1}^n \max_{a_{ij} \neq 0} \{i+1-j\} . \quad (1.5)$$

Note that since it is assumed that $\mathbf{A}$ has a symmetric sparsity pattern, it follows that

$$P = \sum_{i=1}^n f_i . \quad (1.6)$$

In a frontal algorithm, the average number of floating-point operations in a single elimination step is proportional to the r.m.s. wavefront and the maximum amount of storage required for the frontal matrix during the Gaussian elimination is dependent upon the maximum wavefront. Moreover, the total storage required and the amount of work involved in the back substitution stage depends on the profile of the matrix. Thus the elements should be numbered in such a way as to reduce F , $\tilde{F}$, and P .

In recent years, many algorithms for automatically ordering finite elements have been proposed in the literature. These include the methods described by Akin and Pardue (1975), Bykat (1977), Razzaque (1980), Pina (1981), Sloan and Randolph (1983), Fenves and Law (1983), Sloan (1986), Shephard, Baehmann, and Giece (1988), Duff, Reid, and Scott (1989), Kaveh (1991), Koo and Lee (1992), Medeiros, Pimenta, and Goldenberg (1993), and Paulino, Menezes, Gattass, and Mukherjee (1994). Each of these algorithms is designed to reorder **all** the elements in the finite-element domain (they are global reordering schemes). However, in a multiple front method, the finite-element domain is partitioned into a (small) number of subdomains and a frontal decomposition is performed on each subdomain separately (Duff and Scott 1994a, 1994b). For a multiple front method, a global reordering algorithm is unlikely to provide efficient element orderings.

Consider a finite-element domain which has been partitioned into (non-overlapping) subdomains. Variables which belong to a single subdomain are termed *internal variables* and variables which lie on the interface boundaries of the subdomains are called *interface variables*. In general, each subdomain will contain some elements whose variables are all internal variables and some elements with both internal and interface variables. Elements containing only internal variables are called *internal elements* and those with one or more interface variables are called *interface elements*. Internal variables may be eliminated as soon as they are fully summed (provided, of course, that stability criteria are satisfied) but interface variables cannot be eliminated within a subdomain since they are shared by more than one subdomain. In a multiple front algorithm, a frontal solver is applied to each subdomain separately. At the end of the assembly and elimination processes for the subdomains, for each subdomain i there will remain a subdomain frontal matrix $\mathbf{F}_i$ and a corresponding frontal right-hand side vector $\mathbf{c}_i$ (or matrix, if there are multiple right-hand sides) satisfying

$$\mathbf{F}_i \mathbf{y}_i = \mathbf{c}_i . \quad (1.7)$$

If there are n_{sub} subdomains, n_{sub} equations of the form (1.7) are generated and these may be assembled to give

$$\mathbf{F}\mathbf{y} = \mathbf{c}, \tag{1.8}$$

where the order of the matrix $\mathbf{F}$ is the number of interface variables (plus any internal variables not eliminated for stability reasons). By treating each of the subdomain frontal matrices $\mathbf{F}_i$ as an elemental matrix, the system (1.8) may also be solved by a frontal scheme. Once (1.8) has been solved, the results for the interface variables must be passed back to the subdomains so that back-substitution for the internal variables can be performed (for details, see Duff and Scott 1994a, 1994b).

In the multiple front algorithm, once an interface variable has entered the subdomain frontal matrix it cannot be eliminated. Therefore, an element ordering scheme which yields an efficient global ordering may give very inefficient orderings when applied to a single subdomain. The aim of this study is to develop a new element ordering scheme which is efficient when used with the multiple front algorithm.

The remainder of this paper is organised as follows. In Section 2 we introduce some basic concepts from graph theory. In Section 3 we discuss the main features of the algorithm used by the Harwell Subroutine Library (HSL) (Anon 1993) routine MC43, which is designed to resequence elements for use with the HSL frontal code MA42 (Duff and Scott 1993, 1995), and look at why this algorithm is unsuitable for reordering elements for use with the multiple front algorithm. In Section 4 we propose two new schemes for ordering elements in a subdomain and in Section 5 we report on the performance of these methods when used on a range of practical problems. Finally, in Section 6 concluding remarks and comments are made.

2 Graphs and finite-element domains

Associating graphs with finite elements is useful when developing element reordering algorithms. In this section, we briefly recall some basic definitions from elementary graph theory which are of relevance to this paper.

An *undirected graph* G is defined to be a pair (V, E) , where V is a finite set of *nodes* (or *vertices*), and E is a finite set of *edges* defined as unordered pairs of distinct nodes. A *labelling* (or *ordering*) of a graph $G = (V, E)$ with n nodes is a bijection of $\{1, 2, \dots, n\}$ onto V . The integer i ($1 \leq i \leq n$) assigned to a node in V by a labelling is called the *label* (or *number*) of that node. Two nodes i and j in G are said to be *adjacent* if $(i, j) \in E$. The *degree* of a node $i \in G$ is defined to be the number of nodes in G which are adjacent to i , and the *adjacency list* for i is the list of these adjacent nodes. A *path of length k* in G is an ordered set of distinct nodes $(i_0, i_1, \dots, i_k)$ where $(i_{j-1}, i_j) \in E$ for $1 \leq j \leq k$. Two nodes are *connected* if there is a path joining them. A graph G is *connected* if each pair of distinct nodes is connected. Otherwise, G is disconnected and consists of two or more *components*.

The *distance* between nodes i and j in a connected graph G (or in a component of a

disconnected graph) is denoted by $d(i, j)$, and is defined to be the number of edges on the shortest path connecting them. The *diameter* $D(G)$ of G is the maximum distance between any pair of nodes. That is,

$$D(G) = \max \{d(i, j) \mid i, j \in V\}.$$

Nodes at opposite ends of a diameter of G are known as *peripheral* nodes. A *pseudo-diameter* $\delta(G)$ is defined by any pair of nodes i and j in G for which $d(i, j)$ is close to $D(G)$. A pseudo-diameter may be slightly less than, or equal to, the true diameter and is found by some heuristic algorithm. Nodes defining a pseudo-diameter are termed *pseudo-peripheral* nodes.

A *level structure rooted at a node* $r \in V$ is defined as the partitioning of V into levels $l_1, l_2, \dots, l_{h(r)}$ such that

- (i) $l_1 = \{r\}$
- (ii) for $i > 1$, l_i is the set of all nodes that are adjacent to nodes in l_{i-1} but are not in $l_1, l_2, \dots, l_{i-1}$.

The level structure rooted at node r may be expressed as the set $L(r) = \{l_1, l_2, \dots, l_{h(r)}\}$, where $h(r)$ is the total number of levels and is termed the *depth*. The number of nodes on level i is the *width of level i* and is denoted by $|l_i|$. The *width* of the level structure $L(r)$ is given by

$$w(r) = \max_{1 \leq i \leq h(r)} \{|l_i|\}.$$

A *list* is an ordered set. A *priority queue* is a list from which deletions or extractions are made on the basis of a priority function.

A *finite-element domain* is a collection of finite elements in which the elements are joined at their common boundaries and vertices. Finite-element nodes may lie at vertices, along the sides, on the faces, or within the element itself. Associated with each finite-element node is a set of one or more variables corresponding to the freedoms at that node. A convenient way of associating a graph with a finite-element domain consists of choosing the nodes of the graph to be the finite elements and using the interconnection of the finite elements to determine the edges. Relabelling the nodes of this element connectivity graph is equivalent to reordering the elements of the associated finite-element domain and an algorithm which does this is termed a *direct* element reordering algorithm (Duff, Reid, and Scott 1989).

There are several possible ways to use the interconnection of the finite elements to determine the edges of an element connectivity graph. Bykat (1977) defines two elements to be adjacent to one another whenever they share a common edge and describes his algorithm in detail for planar triangular elements. Fenves and Law (1983) generalize this definition to problems in n dimensions, $n = 1, 2, 3$. They define two elements in n dimensions to be adjacent whenever they possess a common boundary of $(n - 1)$ dimensions. Thus in three dimensions two volumetric elements are adjacent if they share a common two-dimensional boundary face; in two dimensions planar elements are interconnected by one-dimensional boundary lines;

and one-dimensional finite elements are adjacent if they have a common finite-element node. Fenves and Law noted that the adjacency of elements cannot always be completely represented by this definition of adjacent elements, since n -dimensional elements are not necessarily connected through $(n-1)$ -dimensional boundaries. In addition, adjacent finite elements do not necessarily have the same geometric dimensions. In such examples, the element connectivity graph may become disconnected, and each component must be numbered independently. This contributes to the difficulties associated with attempting to use this definition of element adjacency.

To avoid these difficulties, Duff, Reid, and Scott (1989) adopt a simpler definition: they define two elements to be adjacent to each other whenever they have one or more variables in common. Using this definition, it is not difficult to generate the associated element connectivity graph; the user does not need to provide information on the different types of elements in the grid other than a list of the variables associated with each finite element. From their numerical experiments, Duff *et. al.* report that using this definition to generate the element connectivity graph and then employing their direct element reordering algorithm did not, in general, lead to a significant increase in the computed maximum and r.m.s. wavefronts compared with those obtained using the element connectivity graph resulting from the adjacency definition of Fenves and Law. The definition of adjacency introduced by Duff *et. al.* has recently been used by Paulino, Menezes, Gattass, and Mukherjee (1994), who term the resulting connectivity graph the *element communication graph*. Throughout the remainder of this paper we will use the element communication (*EC*) graph.

3 Element reordering using MC43

In this section we look at the key features of the direct element reordering algorithm used by MC43 and illustrate the shortcomings of the algorithm if it is used in conjunction with the multiple front method. The MC43 element reordering algorithm exploits the profile reduction algorithm of Sloan (1986). It has three distinct steps:

- (1) selection of a pair of pseudo-peripheral elements (nodes)
- (2) element relabelling
- (3) computation of the maximum wavefront.

In the first step, for each component of the element communication (*EC*) graph, a pair of pseudo-peripheral elements is located. It has been found that, because these elements tend to yield rooted level structures which are deep and narrow, they make good starting elements for profile and wavefront reduction algorithms (see Gibbs 1976 and Sloan and Randolph 1983). The procedure used in MC43 to locate pseudo-peripheral elements is a modification of that described by Gibbs, Poole, and Stockmeyer (1976) and George and Liu (1979). A starting element $s \in G$ of minimum degree is chosen, and the rooted level structure $L(s)$ is generated.

The last level set $l_{h(s)}$ is ‘shrunk’ by retaining only elements with distinct degrees, ties being broken arbitrarily. The level structures rooted at each element in this reduced set (selected in order of increasing degree) are then computed. If, for some $r \in l_{h(s)}$, the depth of $L(r)$ exceeds $h(s)$, r replaces s as the starting element, and the procedure is repeated. If no such element r is found, and e is the element in $l_{h(s)}$ whose associated level structure has the smallest width, the elements s and e are chosen as pseudo-peripheral elements. A ‘short circuiting’ strategy by which wide level structures are rejected as soon as they are detected is incorporated. This algorithm for locating s and e has also recently been used by Medeiros, Pimenta, and Goldenberg (1993).

In the second step of the algorithm, the elements in each component of the EC graph are renumbered to obtain a profile which is smaller than that given by the original labelling of the graph. The pseudo-peripheral elements s and e serve as starting and end elements for the relabelling within their component. The rooted level structure $L(e)$ is generated and the distance $d(e,i)$ of each element i from the end element e is computed. The starting element s is relabelled as element one and a list of elements that are eligible to receive the next label is formed. At each stage in the relabelling process the list of eligible elements comprises those elements which are either adjacent to a element which has been relabelled or are adjacent to a element which is itself adjacent to a relabelled element. The next element to be given a new number is the element among all eligible elements with the highest priority, where the priority P_i of element i is defined to be

$$P_i = -W_1 * ngain(i) + W_2 * d(e,i) - W_3 * nadj(i). \quad (3.1)$$

Here W_1 , W_2 , and W_3 are integer weights, $ngain(i)$ is the number of variables element i will introduce into the front less the number that can then be eliminated, and $nadj(i)$ is the number of elements adjacent to element i which have not yet been relabelled. The basic idea of the algorithm is that, during the reordering process, elements which will make only a small increase to the front size (or will decrease the front size) and which are at a large distance from the end element and have a small number of unlabelled neighbours are labelled first. The values assigned to the weights determines the importance of each of these criteria. As a result of numerical experimentation, in MC43 the weights have the values $W_1 = 10$, $W_2 = 5$, and $W_3 = 1$. We remark that although the EC graph does not take into account the number of variables in each element, the priority P_i of element i given by (3.1) does depend on the number of variables it has.

Once all the elements have been renumbered, MC43 computes the maximum wavefront for the new ordering. If this is larger than the maximum wavefront for the initial ordering, the user is warned that no reduction in the maximum wavefront was achieved and the initial ordering is retained. The value of the maximum wavefront returned by MC43 is useful if the HSL frontal solver MA42 is to be employed since it assists the user in choosing the size of the frontal matrix required and the sizes of the files which will hold the **LU** factors of **A**.

Experience has shown that, when used with MA42, the orderings generated by MC43 are efficient (Duff, Reid, and Scott 1989 and Duff and Scott 1993, 1995). There are, however, weaknesses in each step of the MC43 algorithm if it is employed to resequence the elements in a subdomain for the multiple front algorithm. When locating pseudo-peripheral elements, MC43 does not distinguish between internal and interface elements. As a result, MC43 may choose as a starting element an element containing interface variables. In general, this will be a poor choice since the element numbering should start with elements lying away from the interface boundaries so as to delay the introduction of interface variables into the subdomain front for as long as possible. To illustrate this consider an $N \times 2N$ rectangular domain of rectangular elements. Suppose the elements are initially ordered pagewise parallel to the side of length N and the domain is then partitioned into two, with elements 1 to N^2 in subdomain 1 and elements $N^2 + 1$ to $2N$ in subdomain 2. If MC43 is applied to each subdomain (with the i th element in subdomain 2 corresponding to element $N^2 + i$ in the original domain) the initial orderings will be returned. That is, the ordering for subdomain 1 will start at the exterior boundary and work towards the interface boundary while for subdomain 2 the elements along the interface boundary will be reordered first and those on the exterior boundary last. MC43 is not able to distinguish between these two orderings but, in the multiple front algorithm, the ordering for subdomain 2 will be much less efficient than that for subdomain 1. For example, if $N=8$ and there are five variables at the corners, mid-points of the sides, and centre of each element, numerical experimentation shows that in the multiple front algorithm the two subdomains have r.m.s. wavefronts of 90.6 and 167.0, respectively, and require 1.95×10^7 and 5.85×10^7 floating-point operations, respectively.

In the second step of the MC43 algorithm, the reordering is based on the priorities of the elements computed using (3.1). As we have already seen, an element has a high priority if (a) the net gain it makes to the wavefront is small, (b) it is at a large distance from the end element, and (c) it has a small number of neighbours which have not already been relabelled. Since MC43 is a global algorithm, when calculating the net gain to the wavefront, it is assumed that any variable appearing for the last time may be eliminated. For a subdomain this is no longer true since interface variables may not be eliminated within the subdomain. Failure to take this into account can lead to interface elements being assigned a high priority and being relabelled early in the reordering algorithm. Furthermore, MC43 may select as the end element e an element lying a long way from the interface boundary. In this case, $d(e, i)$ will be large if i is an interface element and will again lead to a high priority for interface elements. This unlikely to yield an efficient ordering.

In the final step, MC43 returns the original and new maximum wavefronts to the user. When a frontal scheme is applied to a domain which has not been partitioned, in general the maximum wavefronts will occur after some, but not all, the elements have been assembled. However, when the domain is partitioned into subdomains, the maximum wavefront may

occur after all the elements have been assembled, and the remaining front then comprises the interface variables (plus any internal variables which have not been eliminated for stability reasons). Since MC43 is unable to distinguish between interface and internal nodes, the maximum wavefronts it returns may be significantly smaller than the wavefronts required by the multiple front method. As already noted, if MC43 finds that the maximum wavefront for the ordering it generates is larger than for the original ordering, it will reject the new ordering and retain the original ordering. But in the multiple front algorithm, the rejected ordering can be more efficient than the accepted ordering. To illustrate this we took the test problem AEAC5081 and partitioned the domain into 8 subdomains using the code Chaco (see Section 5 for details of the test problems and Chaco). MC43 was applied to subdomain 5. The maximum wavefront for the initial ordering was 68. MC43 generates an ordering with a maximum wavefront of 75 so MC43 accepts the original ordering. In the multiple front algorithm, the original ordering gives maximum and r.m.s. wavefronts of 172 and 144.2, respectively, while the rejected ordering yields wavefronts of 172 and 133.0. Thus, for this subdomain, the ordering which was rejected by MC43 provides a more efficient ordering than the one which was accepted.

From the above discussion it is clear that the global reordering algorithm implemented by the code MC43 is not suitable for reordering the elements in a subdomain. Other global reordering schemes applied to a subdomain suffer similar problems to those experienced by MC43 because they too have no concept of internal and interface elements. To be able to use a multiple front algorithm effectively we need to develop a reordering scheme which takes account of interface variables; this is the subject of the rest of this paper.

4 Reordering elements in a subdomain

Algorithms to reorder elements in a subdomain require knowledge of the interface variables. The interface variables also need to be known when the frontal code MA42 is used to implement the multiple front algorithm outlined in Section 1. We have already developed an HSL routine, MA52A, which will generate the list of interface variables for a subdomain. Following the terminology of Duff and Scott (1994a, 1994b), this list is termed the *guard element*.

Intuitively, reordering of a subdomain should begin well away from the interface boundary. We consider two approaches for finding a suitable starting element and we compare the performance of the two approaches in Section 5. Throughout Sections 4.1 to 4.4 we will assume that the element connectivity (*EC*) graph is connected; in Section 4.5 we will consider the case of it having more than one component.

4.1 Approach I

In Approach I we locate starting and end elements s and e using a modified version of the algorithm used by MC43. The steps in finding s and e are as follows.

- (1) Generate a list of interface variables. Flag all elements containing interface variables as interface elements. The unflagged elements are internal elements.
- (2) Select an element s of minimum degree (that is, an element with the smallest number of adjacent elements), if necessary breaking ties by giving preference to internal elements.
- (3) Generate the rooted level structure $L(s) = \{l_1, l_2, \dots, l_{h(s)}\}$.
- (4) Sort the elements in the last level $l_{h(s)}$ in ascending order of their degrees.
- (5) Shrink $l_{h(s)}$ to a list $\tilde{l}$ by retaining only one element of each degree, if necessary breaking ties by giving preference given to internal elements.
- (6) Set $w = \infty$
- (7) For each element $r \in \tilde{l}$ in order of increasing degree, generate $L(r)$. If $h(r) > h(s)$ and $w(r) \leq w$, set $s = r$ and **go to 4**. Else if $w(r) \leq w$, set $e = r$ and $w = w(r)$.
- (8) If s is an internal element **go to 9**. Else if s is an interface element and e is an internal element then set $s = e$ and $e = s$ and **go to 9**. Else s and e are interface elements. Generate $L(s)$ and consider the elements in the middle level set $l_{h(s)/2}$. If there are no internal elements **go to 9**. Else choose s to be the internal element in $l_{h(s)/2}$ of minimum degree, breaking ties arbitrarily.
- (9) Accept s and e as starting and end elements, respectively. Compute the distance $d(e, i)$ of each element i in the subdomain from the end element e .

We remark that the choice for s made in step 8 in general ensures that we begin the element renumbering away from an interface boundary. The case when the elements s and e found in step 7 are both interface elements may occur if the subdomain has interface variables on all sides. Choosing the starting element to lie in the middle level set $l_{h(s)/2}$ is then a way of starting the element reordering with an element at the approximate centre of the subdomain.

4.2 Approach II

In our second approach, we treat the guard element g for the subdomain as an extra element. Suppose there are $nelt$ finite elements in the subdomain and let g be labelled element $nelt + 1$. Let the element communication graph for this augmented set of elements be denoted by $EC(g)$. Recall that in an element communication graph, two elements are adjacent if they have a variable in common. Since the variables in the guard element are the interface variables, in the $EC(g)$ graph the guard element is adjacent to all the interface elements. By considering the rooted level structure $L(g)$, we can locate elements which lie away from the interface boundary and which should make a good choice for the starting element for reordering elements in the subdomain.

Using this idea, the steps in Approach II for choosing a starting element s are as follows.

- (1) Generate the guard element g .
- (2) Generate the element communication graph $EC(g)$.
- (3) Generate the rooted level structure $L(g) = \{l_1, l_2, \dots, l_{h(g)}\}$.
- (4) Shrink $l_{h(g)}$ to a list $\tilde{l}$ by retaining only one element of each degree (with breaking ties arbitrarily).
- (5) Set $h = 1$.
- (6) For each element $r \in \tilde{l}$ in order of increasing degree, generate $L(r)$. If $h(r) > h$, set $h = h(r)$, $s = r$.

In this way, the starting element is chosen to be at a maximum distance from the interface boundary. For a subdomain with interface variables on all sides, Approach II provides a straightforward way of locating an element at the approximate centre of the subdomain. An advantage of this approach is that it avoids the need to locate pseudo-peripheral elements and, except when all the elements in the subdomain are interface elements, the starting element will always be an internal element. The end element is taken to be the guard element, that is, $e = g$. The distance $d(e, i)$ is the distance of element i from the interface boundary. A disadvantage of Approach II is that (s, e) may not provide a good estimate to the diameter of the element communication graph. Furthermore, the $EC(g)$ graph has more edges than the EC graph, the number of extra edges being dependent upon the number of interface variables. The $EC(g)$ graph therefore takes longer to generate than the EC graph and can lead to Approach II being slower than Approach I. However, in our experience, the difference in the times for Approaches I and II was insignificant compared with the time taken to solve the underlying finite-element problem using the multiple front algorithm (see Tables 5.2 and 5.5 in Section 5).

4.3 Element relabelling

Approaches I and II are procedures for choosing starting and end elements s and e for a subdomain. The algorithm we use to relabel the remaining elements in the subdomain is based on that used by MC43 but is modified to take into account interface variables. A priority queue is used where now the priority P_i of element i given by

$$P_i = -W_1 * ngain(i) + W_2 * d(e, i) - W_3 * nadj(i) - W_4 * nint(i). \quad (4.1)$$

Here W_1 , W_2 , W_3 , and W_4 are integer weights. The quantity $ngain(i)$ is the number of variables element i will introduce into the front less the number of internal variables that can then be eliminated. Since interface variables cannot be eliminated within a subdomain, if element i contains one or more interface variables, $ngain(i)$ will be higher than for the same element in a single domain problem and the element will therefore have a lower priority. As in (3.1), $nadj(i)$ is the number of elements adjacent to element i which have not yet been relabelled.

The quantity $nint(i)$ is the number of interface variables the element will introduce into the front and is nonzero only if i is an interface element. The aim of (4.1) is to give a high priority to internal elements. For Approach II, in which the end element e is the guard element, $d(e, i)$ will be large if element i is well away from the interface boundary and this will lead to these elements being given a high priority. Moreover, for Approach II, $nint(i)$ will be nonzero only if $d(e, i)$ is equal to 1 and so for this approach we set $W_4 = 0$. On the basis of our numerical experiments, we suggest choosing the weights to have values $W_1 = 12$, $W_2 = 6$, $W_3 = 1$, and (for Approach I only) $W_4 = 2$. These weights were used in the numerical experiments reported on in Section 5. We found that, in general, small changes to these weight values had no significant effect on the quality of the orderings obtained but, if the ratios between the weights were substantially altered, for some of our problems the resulting orderings had significantly larger wavefronts. An example to illustrate this is included in Section 5 (see Table 5.6).

4.4 Selecting the element order

In MC43, the maximum wavefront is computed for the original and new element orderings. If the new ordering has a maximum wavefront which is no smaller than the maximum wavefront of the original ordering, the original ordering is retained. For the subdomain problem, returning the maximum wavefront which takes into account the interface variables is useful since the frontal solver will need a frontal matrix of at least this size. However, it is not necessarily the appropriate criteria to use to choose between the original and new orderings. As we have already mentioned, for a subdomain, the maximum wavefront may occur after all the elements have been assembled and, in our experience, the reordering algorithms often failed to produce an ordering with a maximum wavefront which was smaller than that for the original ordering. An alternative criteria is to use the r.m.s. wavefront so that if the new ordering provides a smaller r.m.s. wavefront than the original ordering, the new ordering is chosen.

In general, we have found that the choice we have made between the original user-supplied ordering and the new element orderings based on the r.m.s. wavefront has been the correct one. That is, when the frontal solver has been applied to the subdomain, the number of floating-point operations required by the accepted ordering has been fewer than the number required by the rejected ordering. However, even if stability considerations do not cause pivots to be delayed, it is possible that the accepted ordering may still result in the multiple frontal solver requiring more floating-point operations than would be needed by the rejected ordering. When computing the maximum and r.m.s. wavefronts for a given ordering, it is straightforward to also compute the number of floating-point operations the Harwell frontal solver MA42 will require when applied to the subdomain (assuming stability considerations do not cause variables to remain in the front once they are fully summed). In our numerical experiments we compute the number of floating-point operations and choose the ordering for which this is smallest. In this way, when combined with the multiple front algorithm, the element ordering we use never requires more work than the original ordering.

4.5 Coping with more than one component

For the single domain problem, if the *EC* graph has more than one component, each component is reordered separately. The maximum wavefront and the r.m.s. wavefront are independent of the order in which the components are renumbered. However, for the subdomain problem we need to consider the case of more than one component more carefully since the quality of the element ordering will depend upon the order in which the components are renumbered. To illustrate this consider, for example, a subdomain with an associated *EC* graph with two components, *a* and *b*. Suppose the elements in component *a* are passed to the frontal solver first. Let $\mathbf{F}_a$ denote the frontal matrix after the last element in component *a* has been assembled. $\mathbf{F}_a$ is of order at least $nsub_a$, where $nsub_a$ is the number of interface variables lying within component *a*. After some of the elements in component *b* have been assembled the frontal matrix will be of the form

$$\mathbf{F} = \begin{pmatrix} \mathbf{F}_a & \mathbf{0} \\ \mathbf{0} & \tilde{\mathbf{F}} \end{pmatrix}. \quad (4.2)$$

The frontal solver MA42 treats the frontal matrix as dense and so is unable to exploit the zero blocks in (4.2). Thus unnecessary work will be performed and, unless $nsub_a$ is small, it is better to split the subdomain into two subdomains, corresponding to the two components *a* and *b*. If the subdomain is not partitioned, we renumber the components in increasing order of the number of interface variables they contain. Note that, if there are any components with no interface elements, these components can be renumbered first using the standard MC43 algorithm.

In Approach II, the *EC(g)* graph is generated. Although the *EC* graph may have more than one component containing interface variables, the *EC(g)* graph has only one such component. We could apply the two-step algorithm described in Sections 4.2 and 4.3 directly to this single component graph but this can lead the reordering ‘jumping’ about between the components (that is, some of the elements in the first component may not have been renumbered before some of the elements in another component are renumbered). As might be expected, numerical experimentation shows that this can result in poor orderings. Instead, when the *EC(g)* graph has more than one component, we modify Approach II so that, when constructing rooted level structures $L(r)$ using the *EC(g)* graph, we only include elements lying in the same component of the *EC* graph as the root *r*. This is equivalent to splitting the guard element into $ncomp$ guard elements, where $ncomp$ is the number of components in the *EC* graph. Each of these $ncomp$ guard elements is a list of the interface variables lying in the corresponding component of the *EC* graph. The improvements this can give may be illustrated using problem LOCK1074 from the Harwell-Boeing sparse matrix collection (Duff, Grimes, and Lewis 1989, 1992). We partitioned the finite-element domain into 4 using the code Chaco (see Section 5 for details) and generated the *EC* graph for each subdomain. The *EC(g)* graphs for subdomains 3 and 4 were found to have 3 and 2 components, respectively. Applying

Approach II applied to the single component $EC(g)$ graphs for these subdomains gave r.m.s. wavefronts of 185.8 and 196.3, respectively. However, using the proposed modification, the r.m.s. wavefronts were reduced to 118.4, and 114.7, respectively.

5 The test problems

In this section, we report the results of using our proposed subdomain element resequencing algorithms on a range of test problems. All our numerical experiments were performed on a CRAY Y-MP8I using single precision arithmetic and all the reported CPU timings are in seconds. In our experiments, we have only used a single processor of the CRAY but since the reordering of the elements in each subdomain are independent, this can be done in parallel and the frontal method can be applied to each subdomain in parallel. Throughout this section, all computed wavefronts allow for the interface variables, that is, they are the wavefronts for the multiple front algorithm. For the MC43 element ordering the wavefronts are computed outside the MC43 code.

We first use a model problem in which the elements are nine-node rectangular elements with nodes at the corners, mid-points of the sides, and centre of the element. A parameter to the element generation routine determines the number of variables per node. We have chosen this parameter to be five in our numerical experiments. The elements are arranged in a rectangular grid of size $4N \times 4N$ and are initially ordered pagewise. The grid is partitioned into k equal subdomains of order $4N \times 2N$ if $k=2$, $2N \times 2N$ if $k=4$, $N \times 2N$ if $k=8$, and $N \times N$ if $k=16$. If k is equal to 2 or 4, it is straightforward to order the elements efficiently within each subdomain. However, if the number of subdomains is 8 or 16, even for this simple problem it is less apparent how to do this. In Tables 5.1a, 5.1b, and 5.1c we present results for MC43, Approach I, and Approach II applied to the 4-, 8-, and 16-subdomain problem, respectively. For each subdomain, the maximum wavefront $\mathbf{F}$ and the r.m.s. wavefront $\tilde{\mathbf{F}}$ for the different element orderings are given and the ordering which gives the smallest r.m.s. wavefront is in bold typeface. The CPU times required by the element-reordering algorithms for 2, 4, 8, and 16 subdomains are shown in Table 5.2. In addition, the number of floating-point operations (flops) and the CPU time taken by the multiple front algorithm are given. The number of floating-point operations include the floating-point operations for the frontal algorithm in each subdomain together with the floating-point operations needed for the interface problem (1.8) using the HSL code MA42.

Table 5.1a A comparison of MC43, Approach I, and Approach II for the 4-subdomain model problem.

Subdomain	MC43		Approach I		Approach II	
	F	$\bar{F}$	F	$\bar{F}$	F	$\bar{F}$
1	505	374.2	505	374.2	505	348.9
2	515	383.5	505	374.7	505	348.9
3	745	603.1	505	376.5	505	348.9
4	750	612.7	505	377.0	505	348.9

Table 5.1b A comparison of MC43, Approach I, and Approach II for the 8 subdomain model problem.

Subdomain	MC43		Approach I		Approach II	
	F	$\bar{F}$	F	$\bar{F}$	F	$\bar{F}$
1	385	258.6	385	258.6	385	258.6
2	625	399.1	625	399.1	625	399.1
3	625	399.1	625	399.1	625	399.1
4	385	267.4	385	258.6	395	267.4
5	505	371.0	405	259.0	385	275.0
6	735	504.2	625	399.6	625	399.1
7	735	504.2	625	399.6	625	399.1
8	510	380.1	385	259.0	385	275.0

From Tables 5.1a – 5.1c we see that the orderings generated by Approaches I and II have smaller r.m.s. wavefronts than the MC43 ordering and, in general, they also have smaller maximum wavefronts. For each subdomain in the 4- and 16-subdomain problems, Approach II generates an element ordering with a smaller r.m.s. wavefront than Approach I. In some subdomains in the 8-subdomain problem, Approach I does slightly better than Approach II. However, the quality of the ordering generated by Approach I appears to be more dependent on the initial ordering than the Approach II ordering is. For example, for the 4-subdomain problem, Approach II produces orderings with the same maximum and r.m.s. wavefronts for each of the subdomains but Approach I gives orderings which have slightly different r.m.s. wavefronts in each subdomain. From Table 5.2 we see that MC43 is cheaper to use than Approaches I and II, but the small extra cost entailed in using one of the new approaches is more than justified by the savings which the resulting element orderings give in the CPU timings for the multiple front algorithm. We also observe the amount of work can be reduced by partitioning the domain into more subdomains. However, as the number of subdomains increases, the number of interface variables increases and the work involved in solving the interface problem also increases and this rapidly dominates the computation costs. To increase efficiency further, the multiple front algorithm can be nested (see, for example, Benner, Montry, and Weigand 1987).

Table 5.1c A comparison of MC43, Approach I, and Approach II for the 16 subdomain model problem.

Subdomain	MC43		Approach I		Approach II	
	F	$\bar{F}$	F	$\bar{F}$	F	$\bar{F}$
1	265	190.8	265	190.8	265	179.2
2	385	262.3	385	262.3	385	262.3
3	385	262.3	385	262.3	385	262.3
4	275	199.7	265	191.6	265	179.2
5	385	306.6	385	262.3	385	262.3
6	500	371.9	500	367.5	500	349.0
7	500	371.9	500	367.5	500	349.0
8	395	315.4	405	302.6	385	262.3
9	385	306.6	385	262.3	385	262.3
10	500	371.9	500	367.5	500	349.0
11	500	371.9	500	367.5	500	349.0
12	395	315.4	405	302.6	385	262.3
13	385	300.9	285	190.8	265	179.2
14	495	365.9	385	262.3	385	262.3
15	495	365.9	385	262.3	385	262.3
16	390	310.1	385	190.8	265	179.2

Table 5.2 CPU times (seconds) for reordering the elements together with floating-point operation counts (flops) and CPU times (seconds) for solving the model problem using the multiple front algorithm.

Number of subdomains	MC43			Approach I			Approach II		
	Reorder CPU time	Total flops	Solve CPU time	Reorder CPU time	Total flops	Solve CPU time	Reorder CPU time	Total flops	Solve CPU time
2	0.50	1.82E+10	82.46	0.63	1.7EE+10	79.45	0.55	1.70E+10	77.68
4	0.51	1.85E+10	83.62	0.66	1.04E+10	51.09	0.56	9.05E+09	45.26
8	0.55	1.21E+10	57.66	0.71	9.07E+09	45.11	0.61	9.26E+09	46.22
16	0.65	8.95E+09	44.63	0.83	7.51E+09	38.49	0.71	7.00E+09	36.10

In addition to the model problem, a range of problems arising from practical applications have been used to test and assess the quality of the element reordering algorithms discussed in this paper. The problems range in size from 360 to 23446 elements. A brief description of each test problem is given in Table 5.3. Problems with origin HB were taken from the widely-used Harwell-Boeing sparse matrix collection (Duff, Grimes, and Lewis 1989, 1992). Those with origin AEAT were supplied by Andrew Cliffe of AEA Technology, Harwell Laboratory; those with origin DNVR came from A.C. Damhaug of Det Norske Veritas Research, Norway; and those with origin RALPAR were supplied by R.F. Fowler of the Rutherford Appleton Laboratory. For the RALPAR problems, only element connectivity information was available (that is, lists of the finite-element nodes belonging to each element

were supplied but not the variables associated with each of the nodes). Thus the results we present for the RALPAR problems may only be regarded as an indication of the relative performance of the different reordering algorithms. They will be a good indication of performance if the number of variables per finite-element node is relatively constant. For the other problems, a list of the unknowns for each element in the finite-element domain was available. These lists did not include the constrained variables lying on boundaries with Dirichlet boundary conditions. Using these lists, the element orderings produced by the reordering algorithms may differ slightly from those which would be obtained if complete lists of the variables associated with each element in the finite-element domain were available.

Table 5.3. The test problems

Problem identifier	Origin	Description	Number of variables	Number of elements
CEGB3024	HB	2D cross-section of a reactor core	2996	551
CEGB3306	HB	Framework problem from structural engineering	3222	791
LOCK2232	HB	Framework model of a launch umbilical tower	2208	944
LOCK3491	HB	2D vehicle model	3416	684
AEAC5081	AEAT	Double glazing problem	5081	800
RFLOW1	AEAT	Flow problem	9731	1715
OPT1	DNVR	Part of a condeep cylinder	15449	977
TRDHEIM	DNVR	Trondheim fjord model	22098	813
TSYL201	DNVR	Part of a condeep cylinder	20685	960
JETN	RALPAR	3D pipe model	548	360
CHAM	RALPAR	Part of an engine cylinder	12834	11070
TUBU	RALPAR	Engine cylinder model	26573	23446

Before we could test the element reordering algorithms it was necessary to partition the underlying finite-element domains into subdomains. This problem has itself been the subject of much research in recent years and a variety of methods have been discussed in the literature. These methods include the greedy algorithm of Farhat (1987), bandwidth minimisation (Malone 1988), Keringhan and Lin methods (Keringhan and Lin 1970), the inertial bisection method (see, for example, Simon 1991), recursive graph partitioning (Williams 1991), spectral partitioning methods (Pothén, Simon, and Liou 1990, and Hendrickson and Leland 1992), and multilevel methods (Hendrickson and Leland 1993a). For the RALPAR problems, a partitioning of the finite-element domain into subdomains was supplied. This partitioning was obtained using the recursive graph bisection algorithm with

the Keringhan and Lin refinement implemented in the Rutherford Appleton Laboratory code RALPAR code (Greenough and Fowler 1994a, 1994b). For the HB and DNVR problems and problem AEAC5081, the finite-element domain was partitioned using the Chaco package from Sandia National Laboratories (Hendrickson and Leland 1993b). Spectral bisection with the Keringhan and Lin refinement was used. For the RFLOW1 problem, a partitioning of the domain was supplied with the problem.

In all our tests on the problems described in Table 5.3, we have chosen to partition the domain into at most 16 subdomains. This is consistent with the number of subdomains we anticipate when using the multiple front algorithm. Our experience has been that, as the number of subdomains increases (and consequently the number of interface variables increases), it becomes increasingly important to take into account the interface variables when reordering the elements.

In Tables 5.4a-5.4d we present results for the different element ordering algorithms for the HB, AEAT, DNVR, and RALPAR test problems, respectively. For each problem, the maximum and r.m.s. wavefronts (F and $\tilde{F}$) are given for each subdomain for the original user-supplied element ordering, the MC43 ordering, and the Approach I and Approach II orderings. The ordering which gives the smallest r.m.s. wavefront is in bold typeface. CPU timings for reordering the elements together with floating-point operation counts and CPU times for solving the test problems using the multiple front algorithm are given in Table 5.5. The RALPAR examples are not included in Table 5.5 since lists of variables associated with the finite-element nodes were not available for these examples. For problems for which only the matrix sparsity pattern was available (the HB problems and problem AEAC5081), values for the matrix entries were generated using the Harwell Subroutine Library random number generator FA04.

Table 5.4a A comparison of the original, MC43, Approach I, and Approach II element orderings for the HB test problems.

Problem	Subdomain	Number of elements	Number of interface variables	Original F	Original $\tilde{F}$	MC43 F	MC43 $\tilde{F}$	Approach I F	Approach I $\tilde{F}$	Approach II F	Approach II $\tilde{F}$
CEGB3024	1	67	74	115	84.1	108	72.0	85	62.3	80	55.6
	2	66	82	95	58.0	104	72.3	95	57.4	92	56.9
	3	65	90	128	95.4	96	70.2	97	69.1	92	67.9
	4	71	42	99	70.8	88	61.4	69	37.9	62	34.4
	5	72	46	104	69.5	72	47.7	78	55.4	63	43.1
	6	74	66	100	73.7	80	54.3	81	52.5	78	43.5
	7	69	112	136	91.3	144	94.2	128	86.2	118	82.0
	8	67	72	106	82.2	90	64.4	88	65.2	80	60.1
CEGB3306	1	199	114	188	133.9	138	109.0	123	83.1	122	54.6
	2	197	120	259	186.4	126	69.4	127	63.0	126	56.3
	3	197	108	212	169.6	114	62.4	114	59.1	114	54.8
	4	198	102	223	167.5	126	83.8	108	58.9	109	53.1
LOCK2232	1	117	150	276	197.2	156	100.2	156	102.3	156	92.3
	2	118	120	294	215.2	144	107.0	138	103.0	126	82.1
	3	118	108	246	191.9	126	92.8	126	83.4	114	74.5
	4	117	90	234	166.9	138	98.7	108	67.6	96	45.6
	5	117	120	210	93.3	138	98.7	138	92.3	126	64.4
	6	117	78	300	209.1	126	91.7	120	88.0	84	59.8
	7	123	48	246	157.4	102	74.2	78	49.8	66	41.8
	8	117	162	300	214.7	186	123.3	168	109.1	168	92.3
LOCK3491	1	148	108	171	129.1	234	183.1	171	129.1	147	113.6
	2	179	161	318	215.2	276	196.7	167	99.0	167	120.6
	3	176	199	428	285.0	297	196.1	232	137.7	230	134.2
	4	181	210	294	219.4	210	114.5	186	111.6	176	125.9

From Tables 5.4a – 5.4d we see that, in general, Approach II provides the element ordering with the smallest r.m.s. wavefront. As a result, the Approach II element orderings are generally the most efficient when combined with the multiple front algorithm (Table 5.5). The only test problem for which Approach II does not give the best results is LOCK3491; for this problem, Approach I gives slightly better results. In most examples, MC43 does provide a better ordering (in terms of both the maximum and r.m.s. wavefronts) than the original ordering, but this is not guaranteed. For example, for subdomains 1 and 2 of problem TRDHEIM, the MC43 ordering has larger maximum and r.m.s. wavefronts than the original ordering and when used with the multiple front algorithm, requires more floating-point operations than the original ordering. As expected, Approach I, which is essentially the MC43 algorithm with modifications to allow for interface variables, almost always produces element orderings which are an improvement on those generated using MC43 but tie-breaking means that an improvement is again not guaranteed.

The sizes of the reductions in the r.m.s. wavefronts and in the number of floating-point operations which are achieved using the reordering algorithms are obviously dependent on the original user-supplied element ordering. It is clear that for some of the test problems, the original ordering was reasonable and the reordering algorithms were only able to produce modest savings. From Table 5.5 we see that for problem AEAC5081, Approaches I and II gave savings in the floating-point operation counts of little more than 10 and 20 per cent, respectively. However, for many of the test problems the user was not able to provide a good initial ordering and for these problems the reordering algorithms gave substantial savings. For example, for problem OPT1, Approaches I and II reduced the floating-point operation count by about 68 and 75 per cent, respectively. We performed some additional experiments in which the initial element ordering was arbitrary (that is, the user-supplied ordering was randomly permuted). The savings achieved using the reordering algorithms were impressive. For an arbitrary element order, the total number of floating-point operations required by the multiple front algorithm for problem TRDHEIM was 2.04×10^{10} , but Algorithm II reduced this number by about 96 per cent to 8.25×10^8 .

Table 5.4b A comparison of the original, MC43, Approach I, and Approach II element orderings for the AEAT test problems.

Problem	Subdomain	Number of elements	Number of interface variables	Original		MC43		Approach I		Approach II	
				F	$\bar{F}$	F	$\bar{F}$	F	$\bar{F}$	F	$\bar{F}$
AEAC5081	1	48	139	149	108.6	149	108.6	149	108.6	142	103.9
	2	50	152	165	119.3	165	119.3	165	119.3	169	112.6
	3	53	76	79	57.3	79	53.6	79	51.2	79	48.0
	4	49	154	164	119.3	164	105.1	164	104.9	164	101.4
	5	49	154	157	129.2	164	104.3	164	104.9	164	102.1
	6	49	133	136	97.5	136	97.5	136	97.5	136	80.6
	7	50	110	120	80.2	120	80.2	120	80.2	120	80.2
	8	49	119	122	95.2	122	95.2	136	90.6	136	88.2
	9	49	119	129	102.5	129	102.5	129	102.5	136	88.8
	10	50	110	141	101.6	141	101.6	134	99.8	127	83.5
	11	50	119	129	104.1	129	104.1	129	94.9	129	80.5
	12	51	90	100	66.7	131	85.7	100	66.7	100	66.4
	13	51	90	129	105.5	121	83.6	121	84.1	100	66.4
	14	54	76	123	97.0	82	51.1	79	48.3	79	44.4
	15	50	138	141	93.6	141	93.6	141	93.6	141	93.6
	16	48	139	149	107.6	149	107.6	149	107.6	142	101.8
RFLOW1	1	740	668	698	434.1	678	469.7	698	434.1	672	387.6
	2	575	1149	1462	925.0	1456	865.1	1456	864.8	1468	742.6
	3	400	783	827	497.3	827	482.9	827	482.5	815	408.1

We remark that we found the Chaco code was well able to partition the finite-element domains into subdomains having an (approximately) equal number of elements. The Chaco

code attempts to minimise the total number of interface variables and this will reduce both the amount of data which must be transferred between processors when the multiple front algorithm is run in parallel and the amount of work needed to solve the interface problem. However, Chaco can produce partitions in which the subdomains all have a very different number of interface variables and this can lead to a wide variation in the subdomain wavefronts and hence to poor load balancing. Consider the test problem TRDHEIM. We partitioned the domain into 8 subdomains using Chaco and found that each subdomain had 101 or 102 elements but the number of interface variables ranged from 66 for subdomain 7 to 258 for subdomain 4 (Table 5.4c). After reordering the elements with Approach II, the number of floating-point operations required by the frontal solver for these two subdomains were 3.33×10^7 and 1.12×10^8 , respectively. If all 8 processors on the CRAY Y-MP8I are used to solve the problem (one for each subdomain), processor 7 completes its work in 0.93 seconds while processor 4 takes 1.30 seconds. Domain partitioning algorithms which attempt to balance the number of interface variables between subdomains are required for the multiple front algorithm and we plan to investigate this in the future.

Table 5.4c A comparison of the original, MC43, Approach I, and Approach II element orderings for the DNVR test problems.

Problem	Subdomain	Number of elements	Number of interface variables	Original		MC43		Approach I		Approach II	
				F	$\bar{F}$	F	$\bar{F}$	F	$\bar{F}$	F	$\bar{F}$
OPT1	1	243	904	1638	1193.1	941	679.7	925	486.3	925	468.4
	2	236	533	1161	783.1	539	295.8	539	297.7	539	288.6
	3	244	590	1362	1088.6	967	679.2	876	676.6	602	355.8
	4	254	889	1479	1157.3	1016	788.2	979	747.7	952	774.8
TRDHEIM	1	101	180	216	148.7	282	180.8	216	148.7	216	148.7
	2	102	216	336	203.4	414	302.2	228	141.1	228	128.7
	3	101	174	306	205.0	324	199.8	210	131.9	294	142.2
	4	101	258	414	276.0	360	244.7	366	192.5	306	169.3
	5	102	192	300	226.0	300	226.0	294	213.9	240	155.0
	6	102	210	246	178.6	246	178.6	246	178.6	246	178.6
	7	102	66	150	84.0	150	84.0	150	81.2	162	78.9
	8	102	108	234	166.4	234	166.4	156	85.9	150	82.3
TSYL201	1	120	543	651	486.4	594	430.0	564	436.2	564	422.6
	2	120	582	633	457.6	708	492.7	603	390.6	603	390.6
	3	120	732	915	670.0	753	515.5	753	505.9	753	497.3
	4	120	582	633	457.6	708	492.7	603	390.6	603	390.6
	5	120	576	759	540.3	597	454.1	597	459.4	597	422.4
	6	120	582	708	624.7	633	529.0	603	390.6	603	390.6
	7	120	732	768	564.3	753	512.7	753	494.7	753	501.3
	8	120	582	633	457.6	708	492.7	603	390.6	603	390.6

Table 5.4d A comparison of the original, MC43, Approach I, and Approach II element orderings for the RALPAR test problems.

Problem	Subdomain	Number of elements	Number of interface variables	Original		MC43		Approach I		Approach II	
				F	$\bar{F}$	F	$\bar{F}$	F	$\bar{F}$	F	$\bar{F}$
JETN	1	90	61	93	60.9	65	46.4	64	45.1	65	43.7
	2	90	61	81	57.7	63	44.5	63	43.6	62	42.1
	3	90	77	114	80.5	95	74.6	77	45.9	80	48.5
	4	90	57	96	68.5	84	58.3	64	45.9	57	38.5
CHAM	1	1383	342	449	370.9	377	277.4	397	286.4	343	258.8
	2	1383	403	640	484.7	419	300.2	418	302.8	410	274.0
	3	1384	395	484	405.1	454	318.1	462	317.1	446	308.0
	4	1384	489	571	440.0	575	390.7	546	371.4	549	367.8
	5	1384	505	565	449.8	585	397.9	547	369.1	512	356.1
	6	1384	373	490	404.7	420	294.2	385	287.6	374	256.4
	7	1384	529	610	502.5	581	385.8	572	392.8	530	368.2
	8	1384	513	595	460.0	548	377.3	577	383.4	514	363.2
TUBU	1	5861	296	707	533.6	540	366.0	540	361.1	536	352.8
	2	5861	430	657	475.9	759	550.0	471	316.7	556	337.8
	3	5862	691	1740	1305.6	1017	788.2	743	598.6	725	540.9
	4	5862	736	1179	977.7	916	732.9	739	450.1	737	385.0

Table 5.5 CPU times (seconds) for reordering the elements together with floating-point operation counts (flops) and CPU times (seconds) for solving the test problems using the multiple front algorithm.

Problem	Original		MC43			Approach I			Approach II		
	Total flops	Solve CPU time	Reorder CPU time	Total flops	Solve CPU time	Reorder CPU time	Total flops	Solve CPU time	Reorder CPU time	Total flops	Solve CPU time
CEGB3024	3.70E+07	0.82	0.07	2.75E+07	0.76	0.08	2.34E+07	0.71	0.08	2.04E+07	0.69
CEGB3306	1.82E+08	1.69	0.07	4.97E+07	0.90	0.08	3.39E+07	0.81	0.08	2.39E+07	0.74
LOCK2232	1.44E+08	1.46	0.10	4.65E+07	0.90	0.11	3.75E+07	0.85	0.12	2.48E+07	0.75
LOCK3491	3.62E+08	3.31	0.09	2.13E+08	2.21	0.10	1.09E+08	1.67	0.10	1.14E+08	1.70
AEAC5081	1.44E+08	2.41	0.14	1.34E+08	2.32	0.18	1.30E+08	2.31	0.17	1.12E+08	2.20
RFLOW1	8.98E+09	54.0	0.14	8.08E+09	51.30	0.18	7.98E+09	50.4	0.21	6.85E+09	44.2
OPT1	3.46E+10	148.7	0.29	1.39E+10	66.27	0.36	1.12E+10	57.8	0.39	8.61E+09	47.9
TRDHEIM	1.24E+09	12.56	0.22	1.40E+09	13.23	0.29	8.17E+08	10.49	0.28	6.95E+08	9.85
TSYL201	1.40E+10	69.54	0.27	1.12E+10	62.12	0.34	9.79E+09	53.52	0.34	9.57E+09	52.19

In Table 5.6 we illustrate the sensitivity of Approach II to the choice of the weights used in

the priority function (4.1). Results are given for test problems TUBU (4 subdomains) and RFLOW1 (3 subdomains). For each subdomain, the smallest r.m.s. wavefront is in bold typeface. These two problems were selected since they illustrate that for some problems the element ordering can be sensitive to the choice of weights while other problems are much less sensitive. We see that the maximum and r.m.s. wavefronts for RFLOW1 are relatively insensitive to the changes in the weights we considered. For problem TUBU, small changes to the recommended values of 12, 6, and 1 for W_1 , W_2 , and W_3 , respectively, have some effect on the wavefronts but the effects are small. However, for this problem, making substantial changes to the ratios between the weights (such as giving the weights equal values) results in significant increases in the wavefronts in some of the subdomains.

Table 5.6 The sensitivity of Approach II to the priority function weights W_1 , W_2 , W_3 .

Weights W_1 W_2 W_3			Subdomain	TUBU		RFLOW1	
				F	$\bar{F}$	F	$\bar{F}$
12	6	1	1	536	352.8	672	387.6
			2	556	337.8	1468	742.6
			3	725	540.9	815	408.1
			4	737	385.0		
10	5	1	1	544	353.5	672	387.7
			2	558	346.9	1468	742.6
			3	721	541.2	815	408.1
			4	744	389.5		
12	6	2	1	538	358.1	672	387.8
			2	449	317.5	1468	743.1
			3	749	551.6	815	408.1
			4	807	400.7		
2	1	2	1	592	446.1	672	395.4
			2	657	475.9	1468	759.5
			3	942	723.1	818	410.3
			4	792	430.7		
1	0	0	1	707	533.6	672	387.5
			2	476	345.0	1468	745.5
			3	822	587.5	815	408.0
			4	825	406.3		
0	1	0	1	707	533.6	698	434.1
			2	657	475.9	1462	925.0
			3	1740	1305.6	827	497.3
			4	1179	977.7		
1	1	1	1	620	437.5	672	394.4
			2	657	475.9	1468	758.6
			3	885	639.9	818	409.4
			4	800	426.8		

6 Conclusions

In this study we have proposed two algorithms for reordering elements for use with a multiple front algorithm. This problem is more complicated than the problem of resequencing elements for a frontal solver on a single domain since it is necessary to distinguish between variables which can be eliminated once they are fully summed and interface variables that cannot be eliminated within the subdomain. As far as the author is aware, this is the first time this problem has been addressed. The two approaches we have considered involve two different ways of locating a suitable starting element s for the reordering procedure. Once a starting element has been selected, both methods use a modification of the method of Sloan (1986) to reorder the remaining elements. Our first method (Approach I) for choosing s is based on finding pseudo-peripheral nodes of the element communication graph. The second method (Approach II), introduces an artificial element, the guard element, and uses this extra element to find an element lying as far from the interface boundary as possible and uses this to start the reordering. We have tested both approaches on a range of problems and compared their performance with that of the HSL code MC43, which is designed for single domain problems. Both approaches give significant improvements over MC43 and Approach II was almost always the method of choice. On the basis of our findings, a code MC53 implementing Approach II has been developed and will be included in the Harwell Subroutine Library. Since we saw from Table 5.6 that the quality of the element ordering can be dependent on the weights used in the priority function, these weights are passed to MC53 as control parameters. They are given the default values of 12, 6, and 1 (which are the values we recommend on the basis of our numerical experiments) but the user may choose to reset one or more of these values.

Availability of the code

The frontal solver MA42 and the element reordering algorithm MC43 are included in Release 11 of the Harwell Subroutine Library (HSL) (Anon 1993). We plan to include the MA52 package, which allows MA42 to be used to implement a multiple front algorithm, in Release 12 of HSL. The subdomain element reordering code MC53 will also be included in Release 12. Anyone interested in using HSL routines should contact the HSL Manager: Ms L Thick, Harwell Subroutine Library, AEA Technology, Building 8.19, Harwell, Oxfordshire, OX11 0RA, England, tel (44) 235 432688, fax (44) 235 432989, or e-mail libby.thick@aea.org.uk, who will provide details of price and conditions of use.

Acknowledgments

I would like to thank Andrew Cliffe, Christian Damhaug, and Ron Fowler for supplying me with some of the test problems used in this study. I am also grateful to Iain Duff of the Rutherford Appleton Laboratory for his interest and to Alex Pothen of Old Dominion University for some interesting email discussions.

References

- Anon (1993). Harwell Subroutine Library. A catalogue of subroutines (Release 11), AEA Technology, Oxfordshire, England.
- Akin, J. E. and Pardue, R. M. (1975). Element resequencing for frontal solutions. In: *Mathematics of Finite Elements and Applications*, ed. J. R. Whiteman. Academic Press.
- Benner, R. E., Montry, G. R., and Weigand, G. G. (1987). Concurrent multifrontal methods: shared memory, cache, and frontwidth issues. *Int. J. Supercomputer Applics.* **1**, 26-44.
- Bykat, A. (1977). A note on an element ordering scheme. *Int. J. Numer. Meth. Engng.* **11**, 194-198.
- Duff, I. S. (1983). Enhancements to the MA32 package for solving sparse unsymmetric matrices. Harwell Report AERE R11009, HMSO, London.
- Duff, I. S., Grimes, R. G. and Lewis, J. G. (1989). Sparse matrix test problems. *ACM Trans. Math. Softw.* **15**, 1-14.
- Duff, I. S., Grimes, R. G. and Lewis, J. G. (1992). User's guide for the Harwell-Boeing sparse matrix collection. Rutherford Appleton Laboratory Report RAL-92-086.
- Duff, I. S., Reid, J. K., and Scott, J. A. (1989). The use of profile algorithms with a frontal code. *Inter. J. Numer. Meth. Engng.*, **28**, 2555-2568.
- Duff, I. S. and Scott, J. A. (1993). MA42 – a new frontal code for solving sparse unsymmetric systems. Rutherford Appleton Laboratory Report RAL-93-064.
- Duff, I. S. and Scott, J. A. (1994a). In: *Proceedings of the Fifth SIAM Conference on Applied Linear Algebra*, ed. John Lewis, SIAM Press, 567-571.
- Duff, I. S. and Scott, J. A. (1994b). The use of multiple fronts in Gaussian elimination. Rutherford Appleton Laboratory Report RAL-94-040.
- Duff, I. S. and Scott, J. A. (1995). The design of a new frontal code for solving sparse unsymmetric systems. *ACM Trans. Math. Softw.*, submitted.
- Farhat, C. (1998). A simple and efficient automatic FEM domain decomposer. *Computers and Structures* **28**, 579-602.
- Fenves, S. J. and Law, K. H. (1983). A two-step approach to finite element ordering. *Int. J. Numer. Meth. Engng.* **19**, 891-911.
- George, A. and Liu, W. H. (1979). An implementation of a pseudoperipheral node finder. *ACM Trans. Math. Softw.* **5**, 284-295.
- Gibbs, N. E. (1976). A hybrid profile reduction algorithm. *ACM Trans. Math. Softw.*, **2**, 378-387.
- Gibbs, N. E., Poole, W. G., Jr., and Stockmeyer, P. K. (1976). An algorithm for reducing the bandwidth and profile of a sparse matrix. *SIAM J. Numer. Anal.* **13**, 236-250.
- Greenough, C. and Fowler, R. F. (1994a). Partitioning methods for unstructured finite element meshes. Rutherford Appleton Laboratory Report RAL-94-091.

- Greenough, C. and Fowler, R.F. (1994b). RALPAR – The RAL mesh partitioning program. Rutherford Appleton Laboratory Report RAL-94-092.
- Hendrickson, B. and Leland, R. (1992). An improved spectral graph partitioning algorithm for mapping parallel computations. Technical Report SAND92-1460, Sandia National Laboratories.
- Hendrickson, B. and Leland, R. (1993a). A multilevel algorithm for partitioning graphs. Technical Report SAND93-1301, Sandia National Laboratories.
- Hendrickson, B. and Leland, R. (1993b). The Chaco user's guide, version 1.0. Technical Report SAND93-2339, Sandia National Laboratories.
- Hood, P. (1976). Frontal solution program for unsymmetric matrices. *Int. J. Numer. Meth. Engng.* **10**, 379-400.
- Irons, B.M. (1970). A frontal solution program for finite-element analysis. *Int. J. Numer. Meth. Engng.* **2**, 5-32.
- Kaveh, A. (1991). A connectivity coordinate system for node and element ordering. *Comput. Strut.* **41**, 1217-1223.
- Keringhan, B. and Lin, S. (1970). An efficient heuristic procedure for partitioning graphs. *Bell System Technical Journal* **29**, 291-307.
- Koo, B.U. and Lee, B.C. (1992). An efficient profile reduction algorithm based on the frontal ordering scheme and graph theory. *Comput. Strut.* **44**, 1339-1347.
- Malone, J.G. (1998). Automatic mesh decomposition and concurrent finite element analysis for hypercube multiprocessor computer. *Comp. Meth. in Applied Mech. Engng* **70**, 20-58.
- Medeiros, S. R. P., Pimenta, P. M., and Goldenberg, P. (1993). An algorithm for profile and wavefront reduction of sparse matrices with a symmetric structure. *Eng. Comput.* **10**, 257-266.
- Paulino, G.H., Menezes, I.F.M., Gattass, M., and Mukherjee, S. (1994). Node and element resequencing using the Laplacian of a finite element graph: part I-general concepts and algorithm. *Int. J. Numer. Meth. Engng.* **37**, 1511-1530.
- Pina, H.L. (1981). An algorithm for frontwidth reduction. *Int. J. Numer. Meth. Engng.* **17**, 1539-1546.
- Pothen, A., Simon, H.D., and Liou, K.P. (1990). Partitioning sparse matrices with eigenvectors of graphs. *SIAM J. Matrix Anal. Appl.* **11**, 430-452.
- Razzaque, A. (1980). Automatic reduction of frontwidth for finite element analysis. *Int. J. Numer. Meth. Engng.* **15**, 1315-1324.
- Shephard, m.S., P.L. Baehmann, and Giece, K.R. (1988). The versatility of automatic mesh generators based on tree structures and advanced geometric constructs. *Commun. Appl. Numer. Methods* **4**, 379-392.
- Simon, H.D. (1991). Partitioning of unstructured problems for parallel processing. IPProc. Conference on Parallel Methods on Large Scale Structural and Physics Applications. Pergammon Press.
- Sloan, S.W. (1986). An algorithm for profile and wavefront reduction of sparse matrices. *Int. J. Numer. Meth. Engng.* **23**, 239-251.

- Sloan, S. W. and Randolph, M. F. (1983). Automatic element reordering for finite-element analysis with frontal schemes. *Int. J. Numer. Meth. Engng.* **19**, 1153-1181.
- Williams, R. D. (1991). Performance of dynamics load balancing algorithms for unstructured mesh calculations. *Concurrency, Practice, and Experience* **3**, 457-482.